Economics 696F, Homework 5

Due Tues, April 4, 2006, by end of day

Consider the probit model:

$$\Pr(y_i = 1|x_i) = \Phi(\alpha + \beta x_i),$$

where y_i and x_i are scalars, and conditional on α, β, x, the y_i are independent.

The data for this exercise are available in ascii format in the file `hw5.dat`. (The first column of `hw5.dat` contains the vector of observations on y_i, and the second column contains the vector of observations on x_i.)

Consider two priors for α, β:

$$p(\alpha, \beta) \propto 1,$$
$$p(\alpha, \beta) \propto \exp(-\frac{100}{2}\beta^2),$$

1. Calculate the marginal posterior $p(\beta|y, x)$ under the two prior specifications. To perform the calculations, discretize the parameter space for α as

   $$\alpha \in \{-10.0, -9.9, \ldots, 9.9, 10.0\}.$$

   Likewise, discretize the parameter space for β as

   $$\beta \in \{0, .01, .02, \ldots, .49, .50\}.$$

   See below for more information on discretization.

   Plot the two marginal posterior distributions on one plot, and discuss their similarities and/or differences.

2. Write a data augmentation algorithm, assuming a uniform prior on both α and β, to simulate the posterior distribution of (α, β). Initialize the parameters at $(0, 0)$ and run the algorithm for 500 iterations. Save the draws for α and β; you do not need to save the draws for any latent variables for the purposes of this exercise.

   Plot α and β over time and discuss whether it appears that they have "converged." Provide an approximate posterior mean and posterior variance for β. Discard the first 100 iterations and plot a histogram of the last 400 draws for β.

Discretization Method

Suppose we have a likelihood $p(z|\theta)$ and a prior $p(\theta)$ on some parameter space Θ.

One possibility is to discretize the parameter space Θ. Suppose that $\Theta = \{\theta_1, \dots, \theta_K\}$. Then we would have

$$p(\theta_j|z) = \frac{p(\theta_j)p(z|\theta_j)}{\sum_{k=1}^{K} p(\theta_k)p(z|\theta_k)}. \qquad j = 1, \dots, K.$$

The sum in the denominator only needs to be done once, and can be implemented on a computer since Θ is finite. Likewise, calculating expectations with respect to the posterior distribution would only involve finite sums, with the $p(\theta_j|z)$ used as weights.